

# Minimally entangled typical thermal states versus matrix product purifications for the simulation of equilibrium states and time evolution

Moritz Binder<sup>1</sup> and Thomas Barthel<sup>2</sup>

<sup>1</sup>*Department of Physics and Arnold Sommerfeld Center for Theoretical Physics,  
Ludwig-Maximilians-Universität München, Theresienstr. 37, 80333 Munich, Germany*

<sup>2</sup>*Laboratoire de Physique Théorique et Modèles Statistiques,  
Université Paris-Sud, CNRS UMR 8626, 91405 Orsay Cedex, France*

(Dated: November 11, 2014)

For the simulation of equilibrium states and finite-temperature response functions of strongly-correlated quantum many-body systems, we compare the efficiencies of two different approaches in the framework of the density matrix renormalization group (DMRG). The first is based on matrix product purifications. The second, more recent one, is based on so-called minimally entangled typical thermal states (METTS). For the latter, we highlight the interplay of statistical and DMRG truncation errors, discuss the use of self-averaging effects, and describe schemes for the computation of response functions. For critical as well as gapped phases of the spin-1/2 XXZ chain and the one-dimensional Bose-Hubbard model, we assess computation costs and accuracies of the two methods at different temperatures. For almost all considered cases, we find that, for the same computation cost, purifications yield more accurate results than METTS – often by orders of magnitude. The METTS algorithm becomes more efficient only for temperatures well below the system’s energy gap. The exponential growth of the computation cost in the evaluation of response functions limits the attainable timescales in both methods and we find that in this regard, METTS do not outperform purifications.

PACS numbers: 05.30.-d, 02.70.-c, 75.10.Pq, 05.30.Jp

## Contents

<b>I. Introduction</b>	<b>1</b>
<b>II. Matrix product purifications</b>	<b>2</b>
<b>III. METTS sampling</b>	<b>3</b>
A. Algorithm for static observables	3
B. Errors: Statistics, truncations, and Trotter	3
C. Exploiting self-averaging	4
<b>IV. Comparison for static correlators</b>	<b>5</b>
A. Procedures to compare efficiencies	5
B. Spin-1/2 XXZ chain	5
C. 1D Bose-Hubbard model	8
<b>V. Algorithms for response functions</b>	<b>8</b>
A. The simple scheme A	9
B. Scheme B is not useful for METTS	9
C. Reaching longer times with scheme C	9
D. Comparison	11
<b>VI. Response in the spin-1/2 XXZ chain</b>	<b>11</b>
<b>VII. Conclusions and discussion</b>	<b>11</b>
A. Truncations in our tDMRG simulations	12
B. Choices for the METTS collapse	13
C. Quantification of computation costs	13
D. Exact versus quasi-exact reference data	14
<b>References</b>	<b>16</b>

## I. INTRODUCTION

Finite-temperature correlation and response functions of quantum many-particle systems are of great interest. They provide insights into the many-body physics and allow to compare theoretical models to experimental results. However, their precise computation remains challenging. For many relevant models, one has to rely on the development of efficient numerical techniques. The most successful method for the study of strongly correlated one-dimensional (1D) systems is the density-matrix renormalization group (DMRG) which is based on matrix product states (MPS) [1–3]. While DMRG was originally designed to study ground states of 1D systems, its extension to the time evolution of quantum states within tDMRG [4–6] allows for the simulation of quenches and response functions. Based on this extension, quite different methods for the simulations at finite temperatures have been developed. One of them rests on a purification [44] of the density matrix [7–9] which can be encoded in matrix product form [10, 11]. First, this was successfully applied to study static finite-temperature properties of quantum spin chains [12, 13]. The combination with real-time tDMRG allows for the precise evaluation of finite-temperature response functions and can be applied to compute spectral functions and to study a variety of experimentally relevant systems [14–22]. Despite their success, simulations based on such purifications are often limited with respect to the reachable inverse temperatures and/or times, due to a growth of entanglement which is accompanied by a corresponding growth of computation costs. The search for comple-

mentary approaches led to an algorithm that avoids the direct encoding of the mixed states: Instead of purifying the density matrix, one can sample cleverly chosen pure states, so-called minimally entangled typical thermal states (METTS) [23, 24]. While they can be efficiently encoded in matrix product form as their entanglement is relatively low, they represent well the thermal properties of the system at hand. The METTS algorithm has been successfully applied to study static properties and quantum quenches at finite temperature [25–27]. However, a thorough analysis of its accuracy and efficiency compared to computations using matrix product purifications of the density matrices was lacking.

In this work, we discuss how self-averaging can be used to moderately reduce statistical errors in the METTS algorithm and introduce schemes for the evaluation of response functions using METTS. We compare the accuracies and computation costs of the METTS and purification approaches for the evaluation of finite-temperature correlation and response functions. We focus on two paradigmatic models of interacting quantum systems, namely the spin-1/2 XXZ chain [28–30] and the 1D Bose-Hubbard model [31, 32] at critical as well as non-critical points of their phase diagrams. In contrast to indications and expectations expressed in the earlier literature [23, 27, 33], for almost all cases considered here, we find that, for the same computation cost, the purification approach yields more accurate results than METTS – often by orders of magnitude. METTS become more efficient only for temperatures well below the energy gap of the system. It would be interesting to investigate further whether other approaches as, for example, computations based on the ground state and a few excited states, which can be determined variationally, could outperform the METTS approach for such very low temperatures. For the comparisons, we always use equal total computation costs for both methods, ignoring that METTS simulations can be parallelized more easily than purification simulations by generating independent Markov chains on different computing nodes. This can be taken into account by keeping in mind that, to reduce the presented METTS errors by an order of magnitude, one needs to increase the number of employed computing nodes by at least a factor of 100.

The article is structured as follows. Sections II and III shortly review the algorithms for computing static observables with purifications and METTS, respectively, and discuss the interplay of statistical and truncation errors for METTS (Sec. IIIB) as well as self-averaging (Sec. IIIC). In section V, we introduce for METTS a simple scheme and two more elaborate schemes for the computation of response functions which are analogous to corresponding schemes based on purifications [14, 16, 17]. The main objective of the paper, the efficiency comparison of METTS and purifications, is presented in section IV for static correlations functions in the spin-1/2 XXZ chain and the Bose-Hubbard model, and in section VI for response functions in the XXZ model. Some

technical issues are described in appendices. We summarize and conclude in section VII.

## II. MATRIX PRODUCT PURIFICATIONS

Let us briefly review how to compute finite-temperature expectation values  $\langle \hat{O} \rangle_\beta = \text{Tr}(\hat{\rho}_\beta \hat{O})/Z_\beta$  using matrix product purifications. Here, we work with the canonical ensemble  $\hat{\rho}_\beta = \exp(-\beta \hat{H})$  and  $Z_\beta = \text{Tr} \hat{\rho}_\beta$ .

A state  $|\mathbf{P}_{\hat{\rho}}\rangle \in \mathcal{H} \otimes \mathcal{H}_{\text{aux}}$  is called a purification of the density matrix  $\hat{\rho}$  on  $\mathcal{H}$  if

$$\text{Tr}_{\text{aux}} |\mathbf{P}_{\hat{\rho}}\rangle \langle \mathbf{P}_{\hat{\rho}}| = \hat{\rho}. \quad (1)$$

Choosing the auxiliary Hilbert space  $\mathcal{H}_{\text{aux}}$  isomorphic to the physical Hilbert space  $\mathcal{H}$ , i.e.,  $\mathcal{H} \simeq \mathcal{H}_{\text{aux}}$ , it is simple to give a purification of the infinite-temperature state  $\hat{\rho}_0 = \mathbb{1}$ . It is

$$|\mathbf{P}_{\hat{\rho}_0}\rangle = \bigotimes_i \left( \sum_{\sigma_i} |\sigma_i\rangle \otimes |\sigma_i\rangle_{\text{aux}} \right), \quad (2)$$

where  $|\sigma_i\rangle$  are orthonormal basis states for lattice site  $i$ , and  $|\sigma_i\rangle_{\text{aux}}$  for the corresponding lattice site of the auxiliary system. For the orthonormal basis  $\{|\boldsymbol{\sigma}\rangle = \bigotimes_i |\sigma_i\rangle\}$  of  $\mathcal{H}$ , let  $|X\rangle \in \mathcal{H} \otimes \mathcal{H}_{\text{aux}}$  denote the vectorization of an operator  $\hat{X}$  on  $\mathcal{H}$  such that

$$|X\rangle \equiv \sum_{\boldsymbol{\sigma}\boldsymbol{\sigma}'} \langle \boldsymbol{\sigma} | \hat{X} | \boldsymbol{\sigma}' \rangle |\boldsymbol{\sigma}\rangle \otimes |\boldsymbol{\sigma}'\rangle_{\text{aux}}. \quad (3)$$

In this notation, we have that  $|\rho_{\beta/2}\rangle \in \mathcal{H} \otimes \mathcal{H}_{\text{aux}}$  is according to equations (1) and (3) a purification of the density matrix  $\hat{\rho}_\beta$ .

Because  $|\rho_0\rangle \equiv |\mathbf{P}_{\hat{\rho}_0}\rangle$ , as given in Eq. (2), is a product state, it can be encoded as an MPS with matrices of bond dimension one (cf. appendix A). With  $|\rho_0\rangle$  as the initial state, one can employ imaginary-time evolution, to obtain purifications  $|\rho_{\beta/2}\rangle$  for finite-temperature states  $\hat{\rho}_\beta$ ,

$$|\rho_{\beta/2}\rangle = (e^{-\beta \hat{H}/2} \otimes \mathbb{1}_{\text{aux}}) |\rho_0\rangle. \quad (4)$$

To this purpose one can employ the time-dependent DMRG algorithm (tDMRG) [5, 6] or the almost identical time-evolved block decimation (TEBD) [4]. Specifics of our simulations are summarized in appendix A. Exploiting that

$$Z_\beta = \langle \rho_{\beta/2} | \rho_{\beta/2} \rangle \quad \text{and} \quad \hat{\rho}_\beta = \text{Tr}_{\text{aux}} |\rho_{\beta/2}\rangle \langle \rho_{\beta/2}|,$$

thermal expectation values can be computed by

$$\langle \hat{O} \rangle_\beta = \frac{1}{Z_\beta} \text{Tr} (e^{-\beta \hat{H}} \hat{O}) = \frac{\langle \rho_{\beta/2} | \hat{O} | \rho_{\beta/2} \rangle}{\langle \rho_{\beta/2} | \rho_{\beta/2} \rangle}, \quad (5)$$

where both physical and auxiliary degrees of freedom are summed over.

### III. METTS SAMPLING

#### A. Algorithm for static observables

The strategy employed in the minimally entangled typical thermal states (METTS) algorithm is to approximate thermal expectation values  $\langle \hat{O} \rangle_\beta$  by sampling pure quantum states that have two favorable properties. They represent well the physical properties of the system for the given temperature, and the entanglement of the states is relatively low which makes DMRG calculations efficient. Expressing the trace for the thermal expectation value using some orthonormal basis  $\{|n\rangle\}$  of product states

$$|n\rangle = \bigotimes_i |n_i\rangle, \quad (6)$$

where  $|n_i\rangle$  are (arbitrary) orthonormal basis states for lattice site  $i$ , we have

$$\langle \hat{O} \rangle_\beta = \frac{1}{Z_\beta} \sum_n \langle n | e^{-\beta \hat{H}} \hat{O} | n \rangle. \quad (7)$$

Defining the METTS  $|\phi_n\rangle$  and their probabilities  $P_n$ ,

$$|\phi_n\rangle := \frac{1}{\sqrt{P_n}} e^{-\beta \hat{H}/2} |n\rangle, \quad P_n := \langle n | e^{-\beta \hat{H}} | n \rangle, \quad (8)$$

the thermal average reads

$$\langle \hat{O} \rangle_\beta = \frac{1}{Z_\beta} \sum_n P_n \langle \phi_n | \hat{O} | \phi_n \rangle. \quad (9)$$

Thus, by sampling the states  $|\phi_n\rangle$  according to the probability distribution  $P_n/Z_\beta$ , we can approximate  $\langle \hat{O} \rangle_\beta$  by averaging over  $\langle \phi_n | \hat{O} | \phi_n \rangle$ . The computation cost of DMRG is directly related to the entanglement of the quantum state (see appendix C). Hence, product states (6) are a natural choice because their entanglement entropy is zero and it remains reasonably low during imaginary-time evolution. The sampling is accomplished efficiently by generating a Markov chain of METTS as illustrated in Fig. 1. An arbitrary initial product state  $|n\rangle$  is evolved in imaginary time to obtain the METTS  $|\phi_n\rangle$ . (See appendix A for details on the tDMRG evolution.) Then, the METTS is collapsed through a projective measurement of the operator  $\sum_{n'} |n'\rangle \langle n'|$ , yielding a new product state  $|n'\rangle$  with probability  $p_{n'n} := |\langle n' | \phi_n \rangle|^2$  from which one subsequently computes  $|\phi_{n'}\rangle$  and so on. The transition probabilities obey detailed balance

$$p_{n'n} P_n = |\langle n' | e^{-\beta \hat{H}/2} | n \rangle|^2 = p_{nn'} P_{n'} \quad (10)$$

such that the desired distribution  $P_n$  is indeed the fixed point of this Markov process.

Note that the projective measurement,  $|\phi_n\rangle \rightarrow |n'\rangle$ , can be carried out sequentially, site by site. Starting at some site  $i$ , we go from  $|\phi_n\rangle$  to  $|n'_i\rangle \langle n'_i|$ .




Figure 1: METTS algorithm for the evaluation of static observables. A product state  $|n\rangle$  is evolved in imaginary time up to  $\tau = \beta/2$  and normalized to obtain the METTS sample  $|\phi_n\rangle$ . A projective measurement with transition probabilities  $p_{n'n} = |\langle n' | \phi_n \rangle|^2$  yields a new product state  $|n'\rangle$  which is again evolved in imaginary time etc. Observables are evaluated by averaging the expectation values obtained from the samples.

$|\phi_n\rangle / \sqrt{\pi(n'_i)}$  with probability  $\pi(n'_i) := \|\langle n'_i | \phi_n \rangle\|^2$ . Measuring subsequently on site  $j$ , we go to  $|n'_i n'_j\rangle \langle n'_i n'_j| \cdot |\phi_n\rangle / \sqrt{\pi(n'_i) \pi(n'_j | n'_i)}$  with probability  $\pi(n'_j | n'_i) := \|\langle n'_i n'_j | \phi_n \rangle\|^2 / \pi(n'_i)$  such that, in the end, we arrive at state  $|n'\rangle$  indeed with probability

$$p_{n'n} = \pi(n'_i) \pi(n'_j | n'_i) \pi(n'_k | n'_i n'_j) \dots = |\langle n' | \phi_n \rangle|^2.$$

Due to this, the projective measurement of MPS  $|\phi_n\rangle$  can be done efficiently in a single sweep through the lattice [23, 24].

In order to ensure ergodicity and reduce autocorrelation times, it is useful to switch between different measurement bases  $\{|n^{(k)}\rangle\}$  during the sampling. Details on this and our corresponding choice are described in appendix B.

#### B. Errors: Statistics, truncations, and Trotter

There are two error sources for the evolution of MPS in the framework of tDMRG as described in appendix A. The first is due to truncations of low-weight terms in the Schmidt decomposition of the wavefunction. This error is well-controlled by the truncation threshold (we call it  $\epsilon_P$  for purifications and  $\epsilon_M$  for METTS) which bounds, in every time step, the two-norm deviation  $\|\psi_{\text{trunc}} - \psi\|$  of the truncated MPS from the exactly evolved state. We implement tDMRG using fourth-order Trotter-Suzuki decompositions of the evolution operators with time steps  $\Delta t$  of size 0.125 for purifications and 0.05 for METTS ( $\hbar = 1$ ). These are the second error source. The resulting errors of order  $\Delta t^5$  can only become relevant for very large times and we made sure that they are never dominant for the presented data. Additionally, the accuracy of the METTS sampling algorithm is influenced by

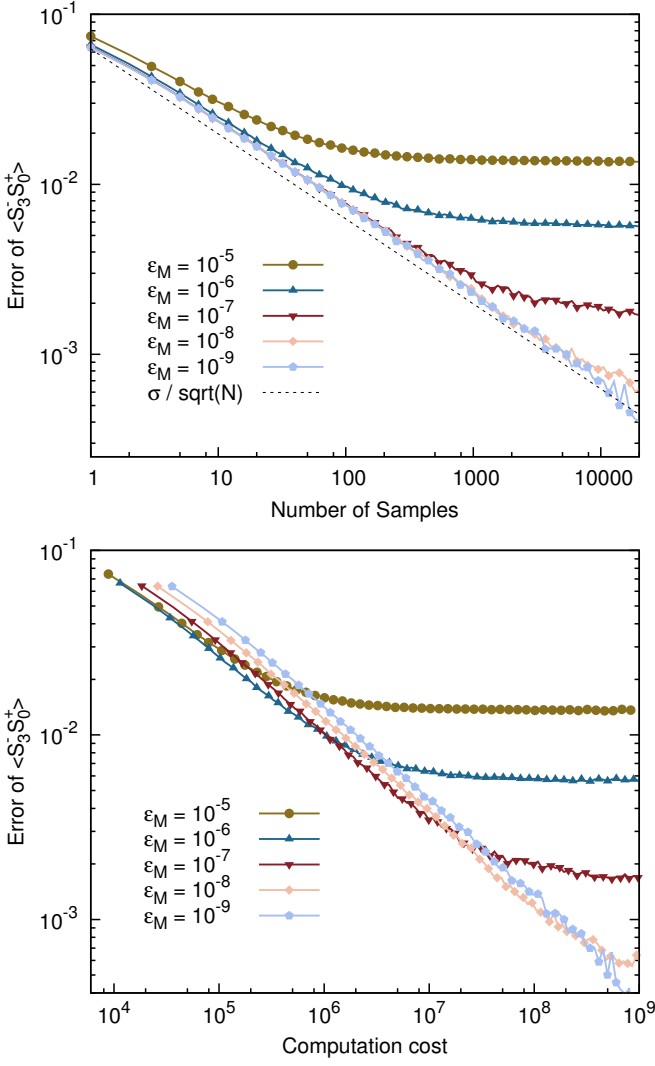


Figure 2: Convergence of the METTS algorithm for the static correlation function  $\langle \hat{S}_3^- \hat{S}_0^+ \rangle_\beta$  in an XX chain [ $\Delta = 0$  in Eq. (11)] of length  $L = 64$  at  $\beta = 4$ . The figure shows the errors for different truncation thresholds  $\epsilon_M$  as a function of the number of samples (top) and as a function of the computation cost (bottom).

a third error source – the statistical error that depends on the number of samples  $N$  used for averaging.

Fig. 2 illustrates the interplay of statistical errors and truncation errors in a METTS computation of the correlator  $\langle \hat{S}_3^- \hat{S}_0^+ \rangle_\beta$  for the exactly solvable 1D XX model [ $\Delta = 0$  in Eq. (11)] at inverse temperature  $\beta = 4$ . The exact solution is shortly described in appendix D. The top panel shows the convergence of the METTS algorithm for different truncation thresholds  $\epsilon_M$  as a function of the number of samples. For low sample numbers, the statistical error dominates and the truncation error is negligible. Autocorrelation times between subsequent samples are short and the statistical error is to a good approximation proportional to  $1/\sqrt{N}$  and independent of the truncation threshold. Once the sample number

reaches a certain  $\epsilon_M$ -dependent threshold, the statistical error has reduced to a magnitude that is comparable to the error induced by the truncations. The curves begin to level out as the relative contribution of the truncation error grows. At a certain point, further samples will not enhance the accuracy of the simulation as the truncations prevent further convergence to the correct value. Of course, the less we truncate (the lower  $\epsilon_M$  is), the longer the  $1/\sqrt{N}$ -convergence persists.

The truncation affects the accuracy in two ways. As every METTS is approximated by a truncated MPS, the values one obtains for each sample are not exact. Additionally, the produced samples will not correspond exactly to the correct probability distribution  $P_n$  of Eq. (9). This is because the transition probabilities depend on the samples and are thus also affected by the truncations.

While lowering the truncation threshold  $\epsilon_M$  yields more accurate results, it also increases the computation cost per sample. In the lower panel of Fig. 2, we present the same errors of the METTS algorithm as in the top panel, but here as a function of the total computation cost which we quantify in an implementation-independent way as described in appendix C. This shifts the simulations with lower truncation thresholds and thus more costly samples to the right. Whereas, for a fixed number of samples (top panel), the accuracy is a monotonic function of  $\epsilon_M$ , this is not necessarily so for fixed total computation cost (lower panel). When plotted against the total computation costs, accuracy curves of METTS simulations with different  $\epsilon_M$  have crossings. For practical simulations, it is therefore important to choose the truncation threshold such that the two error sources are balanced. In the lower panel of Fig. 2, one can easily read off the truncation threshold that is optimal for a given computation cost. While the optimal truncation threshold depends on the specific system studied and the observable that is evaluated, it generally shifts towards lower values of  $\epsilon_M$  with increasing total computation cost.

### C. Exploiting self-averaging

If the considered model is translation invariant, we are free to choose an arbitrary position in the lattice for the evaluation of an observable. While the average of these expectation values will converge to the correct result independent of the position, a single METTS sample itself is not translation invariant. We can thus exploit self-averaging to reduce statistical errors in the METTS algorithm. For a correlation length  $\xi$ , averaging a local observable  $\langle \hat{O}_{x_0} \rangle_\beta$  over a block of  $L'$  sites  $x_0$  corresponds for high and intermediate temperatures to  $\mathcal{O}(L'/\xi)$  statistically independent samples. Therefore, we can expect that the statistical METTS errors reduce by a factor of order  $\sqrt{\xi/L'}$ . For systems with open boundary conditions, one has to restrict the averaging to sites  $x_0$  with a sufficient distance from the boundaries.

We illustrate the effect for the spin-1/2 XXZ chain in



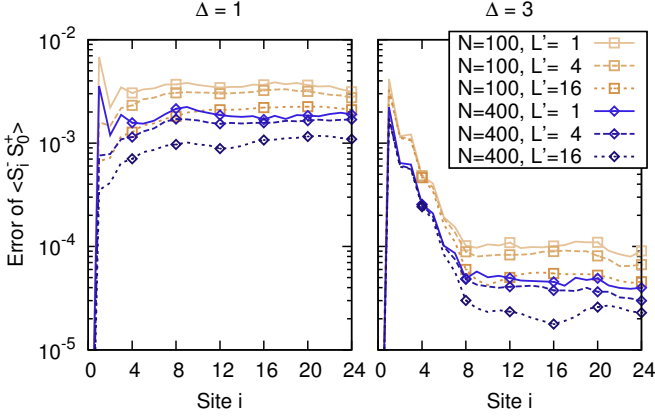


Figure 3: Reducing statistical errors in the METTS algorithm by self-averaging. Shown are errors for the correlator  $\langle \hat{S}_i^- \hat{S}_0^+ \rangle_\beta$  in an XXZ chain (11) of length  $L = 64$  at inverse temperature  $\beta = 4$ . Fixing the truncation threshold to  $\epsilon_M = 10^{-12}$  and the number of samples to  $N = 100, 400$ , the correlator was estimated by averaging over  $L' = 1, 4$  or  $16$  central sites, respectively.

Fig. 3, by averaging the correlator  $\langle \hat{S}_{x_0+i}^- \hat{S}_{x_0}^+ \rangle_\beta$  over different numbers  $L'$  of central sites  $x_0$  and comparing the result to quasi-exact purification data (see appendix D). The observed error reduction is of the expected order of magnitude. For  $\Delta = 3$  and  $\beta = 4$ , the impact of self-averaging is rather small for short distances  $i$ . This is due to the fact that, in this case, the temperature is already well below the gap (cf. Table I) and all excitations occurring in the METTS are of long wavelength. Hence, short-range correlations in the METTS are almost translation invariant.

As the additional computation cost for the spatial averaging of time-local observables (in equilibrium or quench dynamics) is negligible, it is advisable to enhance the METTS accuracy through the self-averaging whenever finite-size effects are well-controlled. Spatial averaging in the evaluation of response function would however require additional real-time evolutions and seems hence not useful.

#### IV. COMPARISON FOR STATIC CORRELATORS

##### A. Procedures to compare efficiencies

In this section we compare the efficiencies of METTS and matrix product purifications, by studying accuracies of static thermal correlation functions for fixed computation costs. To this purpose, accuracies are quantified by the deviations of the obtained expectation values from exact or quasi-exact data as described in appendix D. For the error of  $N$  METTS, we generate several sets of  $N$  subsequent samples, and take the root mean square of the average deviations from the reference

data in each set. The computation cost is quantified in a largely implementation- and platform-independent way as a function of the MPS bond dimensions  $D_i = D_i(\beta, t)$  as detailed in appendix C. For METTS, we average the computation costs of the sample sets.

To take the interplay between statistical and truncation errors (section III B) into account in the assessment of the performances of METTS and purifications, we proceed as follows. We choose a truncation threshold  $\epsilon_P$  for the purification and determine the theoretical computation cost of the simulation. Then we produce METTS samples, using different truncation thresholds  $\epsilon_M$ , and fix the sample set size  $N$  for each  $\epsilon_M$  by the quotient of the purification cost and the average cost per METTS sample such that, for each truncation threshold, the total computation costs of both methods are equal. In the figures, we only present the results for the truncation thresholds  $\epsilon_M$  that yield the best results, i.e., those  $\epsilon_M$  that approximately minimize the error for fixed computation cost, as well as the results for two nearby values of  $\epsilon_M$ .

Setting the total computation costs of both methods equal, ignores that METTS simulations can be parallelized more easily than purification simulations, by generating independent Markov chains on different computing nodes. However, this can be easily taken into account. To reduce the METTS errors, as presented in the following, by an order of magnitude, one has to increase the number of employed computing nodes by at least a factor of 100. This assumes that the extent of thermalization phases at the beginnings of the Markov chains is negligible, and the factor 100 is a lower bound because one also needs to decrease  $\epsilon_M$  (hence, increasing the computation cost per METTS) when the statistical error is being reduced.

##### B. Spin-1/2 XXZ chain

For the spin-1/2 XXZ chain [28–30] with Hamiltonian

$$\hat{H} = \sum_i \left\{ \frac{1}{2} \left( \hat{S}_i^+ \hat{S}_{i+1}^- + \hat{S}_i^- \hat{S}_{i+1}^+ \right) + \Delta \hat{S}_i^z \hat{S}_{i+1}^z \right\}, \quad (11)$$

let us consider three values of the anisotropy parameter: the exactly solvable non-interacting case  $\Delta = 0$  and the isotropic Heisenberg antiferromagnet at  $\Delta = 1$ , which are both critical (gapless), as well as the point  $\Delta = 3$  in the gapped Néel phase. The lattice has size  $L$  and site  $i = 0$  is at the center of the chain. We apply the METTS and purification algorithms, as described in sections II and III A, to compute the static correlation function  $\langle \hat{S}_i^- \hat{S}_0^+ \rangle_\beta$  at inverse temperatures  $\beta = 4$  and  $16$ .

Fig. 4 displays the accuracies of both methods for fixed values of the total computation cost. The columns refer to the three values of the anisotropy parameter  $\Delta$ . The top panels present the absolute value of the quasi-exact correlation function  $\langle \hat{S}_i^- \hat{S}_0^+ \rangle_\beta$ . The panels below display

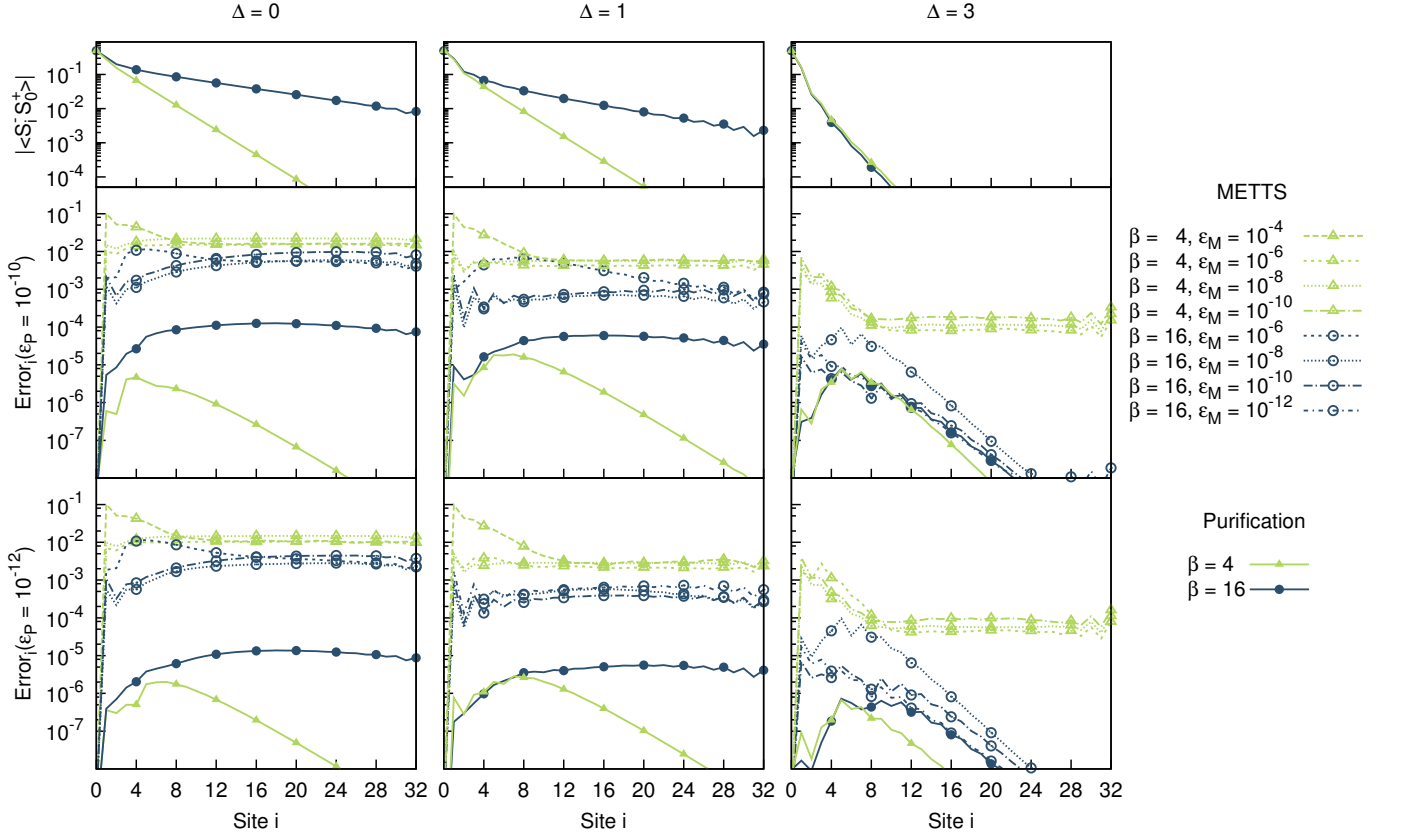


Figure 4: Comparing accuracies of METTS and purifications for the static correlation function  $\langle \hat{S}_i^- \hat{S}_0^+ \rangle_\beta$  in a spin-1/2 XXZ chain (11) of length  $L = 64$  at inverse temperatures  $\beta = 4, 16$ . Site  $i = 0$  is at the center of the chain. The columns refer to the non-interacting case ( $\Delta = 0$ , left), the isotropic Heisenberg antiferromagnet ( $\Delta = 1$ , center) and the gapped Néel phase ( $\Delta = 3$ , right). The top panels show the (quasi-)exact expectation values. Below we compare the errors of the purification with truncation thresholds  $\epsilon_P = 10^{-10}, 10^{-12}$  to several METTS simulations with different truncation thresholds  $\epsilon_M$ . All curves in a given panel that refer to the same temperature share the same computation cost. The corresponding METTS sample set sizes and bond dimensions are specified in Table II.

the errors of the purification simulations with  $\epsilon_P = 10^{-10}$  and  $10^{-12}$  and the errors of several METTS simulations with different truncation thresholds  $\epsilon_M$ . All curves that appear within a panel and refer to the same temperature are based on simulations of equal total computation costs. For each panel, the METTS truncation thresholds  $\epsilon_M$  are chosen such that, for the largest  $\epsilon_M$ , the truncation error dominates; for the lowest, the statistical error dominates; and for the intermediate  $\epsilon_M$ , statistical and truncation errors are balanced in an optimal way such that the error is (approximately) minimized for the given computation cost.

For almost all parameters, the matrix product purification simulations yield more accurate results than the best METTS computations. For  $\beta = 4$ , the errors of the methods differ by up to a few orders of magnitude. When lowering the temperature, entanglement and correlation lengths increase. The increased absolute value of the correlation function is reflected in a correspondingly larger absolute error for the simulations based on purifications. The efficiency of the METTS sampling can increase when lowering the temperatures, especially when it gets suffi-

ciently below the energy gap  $\Delta E$ . In this case, the dimension of the relevant state space to be sampled by METTS is strongly reduced and hence is the statistical error. This is confirmed for the lower temperature  $\beta = 16$  in Fig. 4. For the critical (gapless) systems, the purification approach is still more precise. For the gapped system ( $\Delta = 3$ ), the temperature is with  $\beta = 16$  already substantially below the energy gap,  $\beta \Delta E_{\Delta=3} \sim 10$  (cf. Table I), and METTS can in this case indeed outperform the matrix product purification.

That METTS can become more efficient than the purification is most obvious for the limit  $\beta \rightarrow \infty$ . In this case, every METTS  $|\phi_n\rangle$  for which the initial state  $|\mathbf{n}\rangle$  has nonzero overlap with a ground state will simply be this ground state,  $|\phi_n\rangle = |\text{gs}\rangle$  (up to truncation errors). The purification, on the other hand, evolves to the purification  $|\text{gs}\rangle \otimes |\text{gs}\rangle_{\text{aux}}$  of the ground state density matrix  $|\text{gs}\rangle\langle\text{gs}|$ . As the tensor product of two MPS of bond dimension  $D$  is an MPS with bond dimension  $D^2$ , for the same accuracy,  $\beta \rightarrow \infty$  computation costs for the METTS are reduced by roughly a factor of  $D^3$  [45].

The error of the purification changes significantly as a




Figure 5: Comparing accuracies of METTS and purifications for the static correlation function  $\langle \hat{b}_0^\dagger \hat{b}_i \rangle_\beta$  in the 1D Bose-Hubbard model (12) of length  $L = 32$  at inverse temperatures  $\beta = 4, 16$  with  $\mu = 0.5$  and  $U = 1$ . The local Hilbert spaces were truncated to a maximum of five bosons per site, and site  $i = 0$  is at the center of the chain. The columns refer to the Mott-insulating phase ( $J/U = 0.09375$ , left) with one boson per site in the ground state, the phase boundary ( $J/U = 0.125$ , center) and the superfluid phase ( $J/U = 0.25$ , right). The corresponding METTS sample set sizes and bond dimensions are specified in Table III.

function of the distance  $i$ . Here, the deviation from the (quasi)-exact reference data is mainly due to the truncations. Generally, when we evaluate a correlator based on a truncated matrix product state, the error can grow at short distances, reaching an  $\epsilon_P$ -dependent maximum, before starting to decay as the absolute value of the correlator itself becomes very small. On the other hand,

METTS simulation errors often remain constant at large distances. Even for the exponentially decaying correlation functions, the METTS error stays well above zero. This is clearly a signature of the statistical error. Cases where the METTS error decays with distance, are usually situations where the truncation error dominates over the statistical error. For the computation costs chosen in our study, this is seen for temperatures well below the energy gap, but one can also observe this behavior when choosing very large truncation thresholds.

Spin-1/2 XXZ chain			
	$\Delta = 0$	$\Delta = 1$	$\Delta = 3$
$\Delta E_{L=32}$	0.0475(8)	0.1177(4)	0.6521(9)
$\Delta E_{L=64}$	0.0241(6)	0.0617(6)	0.6242(5)
1D Bose-Hubbard model			
	$J = 0.09375$	$J = 0.125$	$J = 0.25$
$\Delta E_{L=32}$	0.1415(6)	0.032(6)	0.0134(4)
$\Delta E_{L=64}$	0.1395(2)	0.0291(6)	0.0005(0)

Table I: Energy gaps to the first excited state for the spin-1/2 XXZ model (11) and the Bose-Hubbard model (12), determined by DMRG.

Finally, let us shortly discuss an ergodicity issue in the gapped Néel phase. For  $\beta = 16$ , the system is essentially in its ground state. Hence, the weight of the two degenerate Néel states in the thermal state becomes significant. The purified state still obeys the  $Z_2$  symmetry. Every METTS sample is some linear combination of the ground states. Depending on the choice of the collapse basis  $\{|\mathbf{n}\rangle\}$ , severe ergodicity problems can occur in the METTS algorithm. As can be shown, even for the rotation-symmetric random bases that we use (described in appendix B) and which seem to have very good ergodicity properties at first sight, the transition probabilities from one of the degenerate ground states to

the other decay exponentially in the system size  $L$ . Several remedies are available. One can for example avoid this issue by symmetrizing every METTS before measurements. However, in the case of response functions (see section V), this would require an additional real-time evolution and, hence, roughly a doubling of the costs. Alternatively, one could, e.g., use the  $\hat{S}^x$  and the  $\hat{S}^y$  eigenbases for the METTS collapse. Here, we decided to keep the random collapse bases and implicitly average over the two Néel states by using symmetrized observables; for  $\langle \hat{S}_i^-(t) \hat{S}_0^+(0) \rangle_\beta$ , according to

$$\langle \hat{S}_i^-(t) \hat{S}_0^+(0) \rangle_\beta + \langle \hat{S}_i^+(t) \hat{S}_0^-(0) \rangle_\beta = 4 \langle \hat{S}_i^x(t) \hat{S}_0^x(0) \rangle_\beta,$$

such that the measurement is not sensitive to the broken symmetry.

### C. 1D Bose-Hubbard model

The Hamiltonian of the one-dimensional (1D) Bose-Hubbard model [31, 32] is given by

$$\hat{H} = \sum_i \left\{ -J \left( \hat{b}_i^\dagger \hat{b}_{i+1} + \text{h.c.} \right) + \frac{U}{2} \hat{n}_i (\hat{n}_i - 1) - \mu \hat{n}_i \right\} \quad (12)$$

with ladder operators  $\hat{b}_i$  obeying  $[\hat{b}_i, \hat{b}_j^\dagger] = \delta_{ij}$  and the number operators  $\hat{n}_i = \hat{b}_i^\dagger \hat{b}_i$ , hopping  $J$ , onsite interaction  $U$ , and chemical potential  $\mu$ .

We set  $\mu = 0.5$  and  $U = 1$ , and consider three values of the hopping parameter.  $J = 0.09375$  corresponds to the Mott-insulating phase with one particle per site,  $J = 0.125$  places the system close to the phase boundary between the Mott-insulating and the superfluid phase, and  $J = 0.25$  lies in the superfluid phase. In Fig. 5, we show the comparison between METTS and purifications for the Bose-Hubbard model at  $\beta = 4$  and  $\beta = 16$ . Again, all curves in a given panel that refer to the same temperature share the same computation cost. In general, the results are similar to those for the XXZ model. However, the differences between the METTS and purifications errors are larger in the comparison for the Bose-Hubbard model. This can be explained with the smaller energy gaps, compared to those in the spin-1/2 XXZ model. See Table I. For the Bose-Hubbard model at the specified points in the phase diagram, errors of the METTS simulations exceed those of purification simulations by two to three orders of magnitude.

## V. ALGORITHMS FOR RESPONSE FUNCTIONS

The matrix product purification and METTS methods, as described in sections II and III A, can be extended to

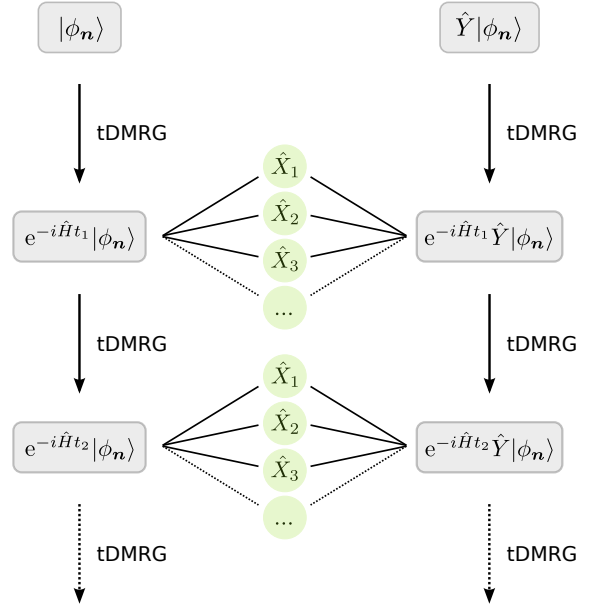


Figure 6: The METTS scheme for the evaluation of response functions used in this work (scheme A). We carry out two independent tDMRG-simulations with the METTS sample  $|\phi_n\rangle$  and the state  $\hat{Y}|\phi_n\rangle$  as initial MPS. At intermediate times  $t_1, t_2, \dots$  we insert operators  $\hat{X}_i$  to evaluate. The values obtained need to be averaged over all METTS samples.

the computation of thermal response functions

$$\langle \hat{X}(t) \hat{Y} \rangle_\beta \equiv \frac{1}{Z_\beta} \text{Tr}(\hat{\rho}_\beta \hat{X}(t) \hat{Y}) \quad (13)$$

with  $\hat{X}(t) \equiv e^{iHt} \hat{X} e^{-iHt}$ . For purifications, several schemes have been suggested and analyzed [14, 16–18]. They have different properties concerning the dependence of computation costs on  $\beta$  and  $t$ . According to the naming introduced in Refs. [17, 18], we address below computation schemes A, B, and C, introduce corresponding schemes for the METTS framework, and discuss their properties.

In Refs. [17, 18], further optimizable classes of schemes have been studied. In comparison to the near-optimal scheme C, they allow to substantially reduce computation costs, for example, for systems with separated energy scales. As there is probably no useful adaption of them for METTS, they will not be addressed here. In related work, Pizorn et al. [34] recently discussed the option of working, for matrix product purifications (equivalently, one can think in terms of matrix product operators [18]), in the Heisenberg picture, i.e., to compute the vectorization  $|X^\dagger(t)\rangle$  of the evolved operator  $\hat{X}^\dagger(t)$  and the vectorization  $|\rho_\beta\rangle$  of the thermal state  $\hat{\rho}_\beta$  to then obtain  $\langle \hat{X}(t) \hat{Y} \rangle_\beta = \langle X^\dagger(t) | \hat{Y} | \rho_\beta \rangle / Z_\beta$ . One may notice that this is just a special case of the class of optimizable schemes suggested in Eq. (9) of Ref. [17] or Eq. (18) of Ref. [18]. We will not discuss it further because this






Figure 7: Computation costs for different evaluation schemes in the METTS algorithm. For the isotropic Heisenberg anti-ferromagnet ( $\Delta = 1$ ) at inverse temperatures  $\beta = 4$  (left) and 16 (right), we show the average cost per sample of the real-time evolution for the evaluation of  $\langle \hat{S}_0^-(t) \hat{S}_0^+(0) \rangle$  in schemes A, B and C (system size  $L = 32$ , truncation threshold  $\epsilon_M = 10^{-10}$ ).

Heisenberg picture scheme is in general computationally suboptimal and has no direct METTS equivalent. Specifically, the cost analysis in Refs. [17, 18] suggests that the Heisenberg picture scheme will typically reach half the maximum times that can be reached by the optimized schemes or the near-optimal scheme C to be described below.

#### A. The simple scheme A

Starting from the matrix product purification  $|\rho_{\beta/2}\rangle$  [Eq. (4)] of the density matrix  $\hat{\rho}_\beta$ , according to evaluation scheme A [14] for the response function (13), we first compute matrix product representations of  $e^{-i\hat{H}t}|\rho_\beta\rangle$  and  $e^{-i\hat{H}t}\hat{Y}|\rho_\beta\rangle$  using tDMRG. The response function is then given by the matrix element

$$\langle \hat{X}(t) \hat{Y} \rangle_\beta = \frac{1}{Z_\beta} [\langle \rho_{\beta/2} | e^{i\hat{H}t} ] \hat{X} [ e^{-i\hat{H}t} \hat{Y} | \rho_{\beta/2} \rangle ]. \quad (14)$$

The vectorization (3) of operators corresponds to the isomorphism between the space  $\mathcal{B}(\mathcal{H})$  of linear maps on the physical Hilbert space  $\mathcal{H}$  and the tensor product space  $\mathcal{H} \otimes \mathcal{H}_{\text{aux}}$ . It allows us to formulate equivalently (and more intuitively) scheme A in terms of matrix product operators (MPOs). Indicating MPOs with square brackets, scheme A reads in this representation simply

$$\langle \hat{X}(t) \hat{Y} \rangle_\beta = \frac{1}{Z_\beta} \text{Tr} \left( [\hat{\rho}_{\beta/2} e^{i\hat{H}t}] \hat{X} [ e^{-i\hat{H}t} \hat{Y} \hat{\rho}_{\beta/2} ] \right). \quad (15)$$

As done in Eqs. (14) and (15), the evolved MPS or MPOs that are used for the evaluation of the response function are always indicated by square brackets in the following.

The METTS equivalent of scheme A is illustrated in Fig. 6. For every sample  $|\phi_n\rangle$ , we compute

$$[\langle \phi_n | e^{i\hat{H}t} ] \hat{X} [ e^{-i\hat{H}t} \hat{Y} | \phi_n \rangle ] \quad (16)$$

using real-time evolution and then average over the values obtained for each sample. In practice this means that we have to carry out two independent tDMRG simulations, using  $|\phi_n\rangle$  and  $\hat{Y}|\phi_n\rangle$  as initial states, up to some maximum time. At any intermediate time-point that we are interested in, we can evaluate the response function  $\langle \hat{X}(t) \hat{Y} \rangle_\beta$  – even for a set of operators  $\{\hat{X}_i\}$  if we wish.

#### B. Scheme B is not useful for METTS

In the context of matrix product purifications, we have the alternative scheme B [16] which reads in the MPO representation

$$\langle \hat{X}(t) \hat{Y} \rangle_\beta = \frac{1}{Z_\beta} \text{Tr} \left( [\hat{\rho}_{\beta/2}] \hat{X} [ e^{-i\hat{H}t} \hat{Y} e^{i\hat{H}t} \hat{\rho}_{\beta/2} ] \right). \quad (17)$$

In comparison to scheme A (15), it corresponds to shifting  $\exp(i\hat{H}t)$  from the first to the second MPO. In Ref. [18], it was explained why this scheme has some advantages at higher temperatures and disadvantages at lower temperatures. The METTS equivalent would be to compute

$$[\langle \phi_n | \hat{X} [ e^{-i\hat{H}t} \hat{Y} e^{i\hat{H}t} | \phi_n \rangle ]. \quad (18)$$

As exemplified in Fig. 7 and explained in the following, its computation costs are unfortunately strictly higher than those of scheme A (16). Both schemes share the cost for computing  $[e^{i\hat{H}t} | \phi_n \rangle]$ . In scheme A, the evolution of  $[e^{-i\hat{H}t'} \hat{Y} | \phi_n \rangle]$  for  $t' = 0 \rightarrow t$  has usually about the same cost, and the entanglement in both states increases typically linearly with  $t'$  everywhere in the system, reaching some value  $\mathcal{E}_t$ . In scheme B (18), the required evolution of  $[e^{-i\hat{H}t'} \hat{Y} e^{i\hat{H}t} | \phi_n \rangle]$  for  $t' = 0 \rightarrow t$  starts with the high entanglement value  $\mathcal{E}_t$ . Due to quasi-locality [35, 36], the entanglement will then reduce in regions of the lattice that are at sufficient distance from the spatial support of operator  $\hat{Y}$ . It will however remain high or even increase in the vicinity of  $\hat{Y}$ .

#### C. Reaching longer times with scheme C

Among the classes of optimizable evaluation schemes studied in Refs. [17, 18], the near-optimal scheme C was found to be very useful and reaches about twice the maximum times reachable with schemes A and B. In the MPO

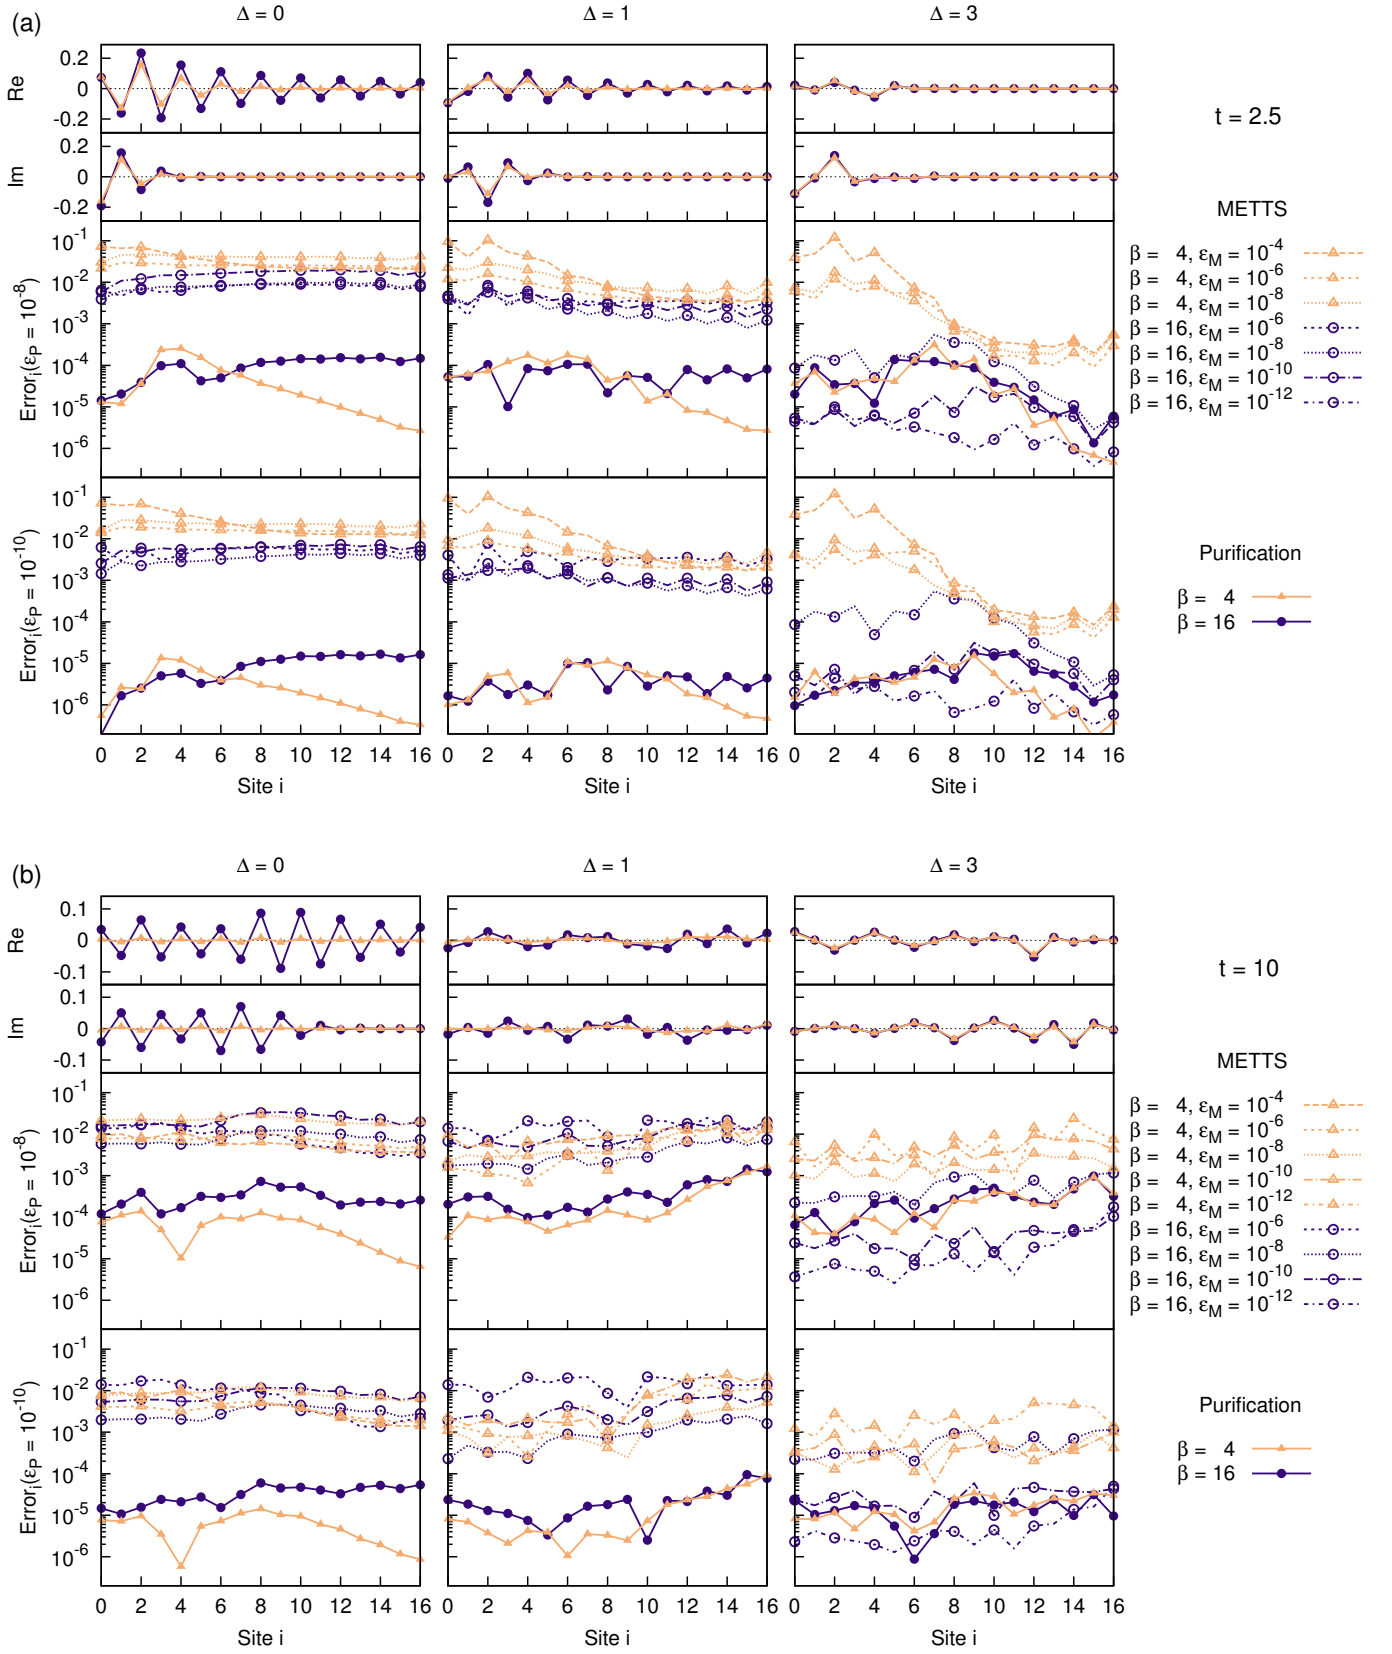


Figure 8: Comparing METTS and purifications for the finite-temperature response function  $\langle \hat{S}_i^-(t) \hat{S}_0^+(0) \rangle_\beta$  in the spin-1/2 XXZ chain at (a)  $t = 2.5$  and (b)  $t = 10$ . In the top panels, we show the (quasi-)exact real and imaginary part of the response function. Below we compare the errors as in Fig. 4. The system size was set to  $L = 32$  and evaluation scheme A was used.

representation, it reads

$$\langle \hat{X}(t)\hat{Y} \rangle_\beta = \frac{1}{Z_\beta} \text{Tr} \left( [\hat{\rho}_{\beta/2} e^{i\hat{H}t/2} \hat{X} e^{-i\hat{H}t/2}] \times [e^{-i\hat{H}t/2} \hat{Y} e^{i\hat{H}t/2} \hat{\rho}_{\beta/2}] \right). \quad (19)$$

In a similar way, we can compute the response function with METTS, by averaging over

$$[\langle \phi_{\mathbf{n}} | e^{i\hat{H}t/2} \hat{X} e^{-i\hat{H}t/2} ] [e^{-i\hat{H}t/2} \hat{Y} e^{i\hat{H}t/2} | \phi_{\mathbf{n}} \rangle] \quad (20)$$

which makes it possible to reach, for the same computational resources, times approximately twice as big as in scheme A.

#### D. Comparison

Due to the cyclic property of the trace and the (trivial) fact that the evolution operators commute with density matrices  $\hat{\rho}_\beta$ , the results of all three schemes converge up to differences in truncation and Trotter errors (see appendix A) to the same thermal average (13). As the truncations in the real-time evolutions are well controlled, the errors of the simulations based on different schemes do not differ significantly if all other parameters are left unchanged. Statistical errors can however differ somewhat [46]. Fig. 7 compares the computation costs of the different evaluation schemes in the METTS algorithm for an autocorrelation function in the isotropic Heisenberg antiferromagnet [ $\Delta = 1$  in Eq. (11)]. While scheme B proves to be the least efficient for the parameters considered here, scheme C indeed reduces the computation costs significantly and thus allows for the evaluation of longer maximum times with METTS.

Nevertheless, we choose to use scheme A in this work because it allows to evaluate response functions  $\langle \hat{X}_i(t)\hat{Y} \rangle_\beta$  for a whole set of operators  $\{\hat{X}_i\}$  and all times  $t$  (up to the maximum reachable time) with only two tDMRG runs per METTS sample. This is very useful for the analysis of response functions like  $\langle \hat{S}_i^-(t)\hat{S}_0^+(0) \rangle_\beta$ , studied in section VI. In contrast, scheme C requires, for every METTS sample, separate tDMRG runs for every required time  $t$  and operator  $\hat{X}_i$ . This is also a drawback in comparison to scheme C for purifications [Eq. (19)], where one obtains results for all times  $t$  with only two tDMRG runs.

#### VI. RESPONSE IN THE SPIN-1/2 XXZ CHAIN

Fig. 8 compares the accuracies of the response function  $\langle \hat{S}_i^-(t)\hat{S}_0^+(0) \rangle_\beta$  in the XXZ chain as computed via METTS and purification simulations for times  $t = 2.5$  and 10. The structure of the plots is as in Fig. 4, i.e. the columns again refer to the three values of the anisotropy parameter  $\Delta = 0, 1, 3$  already considered in section IV B

and all curves that appear within a panel and refer to the same temperature are based on simulations with the same total computation costs. As the thermal response functions are complex-valued, we show both the real and the imaginary part of the (quasi-)exact results.

Generally, the error curves closely resemble the results we obtained for the computation of static correlators in Fig. 4. The purifications are more efficient than METTS except for the gapped system ( $\Delta = 3$ ) at the low temperature  $\beta = 16$ .

Fig. 9 displays the errors and computation costs of both methods as a function of time  $t$  in the evaluation of response functions for  $\Delta = 0$  at inverse temperature  $\beta = 16$ . Each column refers to a truncation threshold  $\epsilon_M$  for the METTS algorithm and shows the errors (top) and computation costs (bottom) for different numbers of samples. The purification errors and costs of three different truncation thresholds  $\epsilon_P$  are shown in each column.

The convergence of the METTS errors is clearly reminding of the behavior shown in Fig. 2. As a function of the number of samples, the error is approximately proportional to  $1/\sqrt{N}$ , until it saturates to an  $\epsilon_M$ -dependent threshold that can be lowered by decreasing the truncation threshold.

The exponentially growing computation costs in both the METTS approach as well as the matrix product purification approach limit the reachable maximum times, i.e., at least qualitatively, METTS seem to have no favorable properties in this respect.

#### VII. CONCLUSIONS AND DISCUSSION

We have studied properties of the METTS algorithm. There is an interplay of statistical and DMRG truncation errors that one should take into account for efficient simulations. While the optimal truncation threshold  $\epsilon_M$  for METTS depends on the specific system studied and the observable that is evaluated, it generally shifts towards lower values of  $\epsilon_M$  with increasing total computation cost. As demonstrated, one can also exploit self-averaging of (approximately) translation invariant systems to reduce statistical errors in METTS simulations for static observables.

We have presented a simple scheme for the evaluation of response functions using METTS and two more elaborate schemes, one of which gives access to longer maximum times but needs a separate simulation for every required point in time.

For spin-1/2 XXZ chains and the 1D Bose-Hubbard model, we have compared the accuracies and computation costs of the METTS and purification approaches for the evaluation of finite-temperature correlation and response functions. For almost all cases considered here, we found in contrast to indications and expectations expressed in the earlier literature that, for the same total computation cost, the purification approach yields more accurate results than METTS – often by orders of mag-

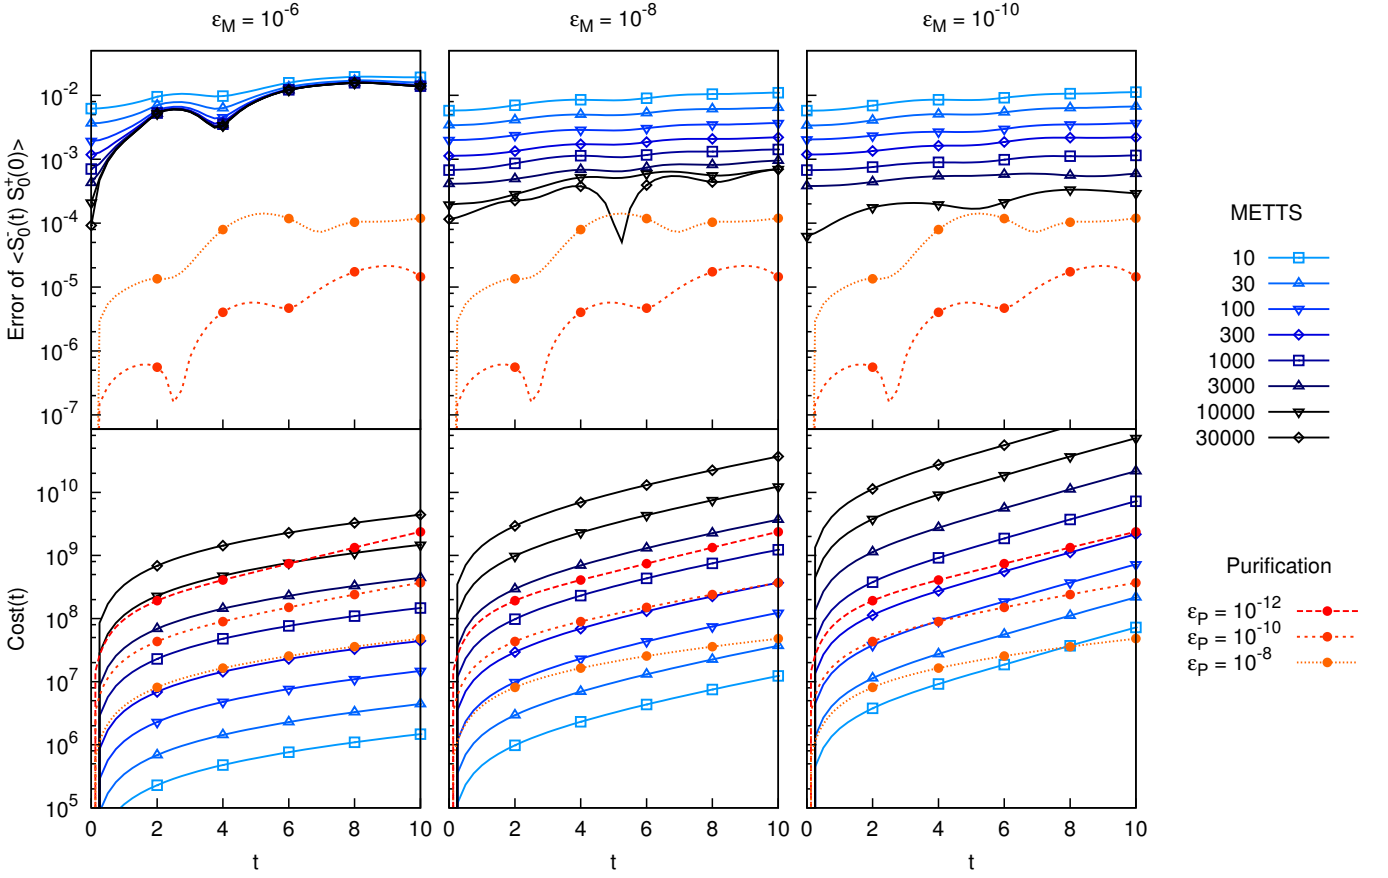


Figure 9: Time-dependence of errors and computation costs for the spin-1/2 XXZ chain of length  $L = 32$  at  $\Delta = 0$  and  $\beta = 16$ . The columns refer to different truncation thresholds in the METTS algorithm. The top panels show the errors for the calculation of  $\langle \hat{S}_0^-(t) \hat{S}_0^+(0) \rangle_\beta$  based on different sample numbers in the METTS algorithm and different truncation thresholds  $\epsilon_P$  in the purification. In the lower panels, the corresponding computation costs are presented.

nitude. METTS become more efficient only for temperatures well below the energy gap of the system ( $\beta \Delta E \gtrsim 10$  in our case). For both methods, we have discussed the temperature dependence of the accuracies, and for correlators also their distance dependence.

It would be interesting to investigate further whether other DMRG approaches for the low-temperature regime, such as computations based on the ground state and a few excited states or a sampling that is restricted to the complement of the ground state space, could outperform the METTS in these cases.

### Acknowledgments

We gratefully acknowledge discussions with U. Schollwöck, S. R. White, and M. Zvonarev.

### Appendix A: Truncations in our tDMRG simulations

The results for the observables and computation costs presented here are essentially independent of the cho-

sen time-evolution algorithm. We use time-dependent DMRG (tDMRG) [4–6] and employ a Trotter-Suzuki decomposition [37, 38] of fourth order, i.e., approximate the time-evolution operator for a (real or imaginary) time step  $\Delta\tau$  according to

$$e^{-\Delta\tau \hat{H}} = \prod_k e^{-a_k \Delta\tau \hat{H}_{\text{even}}} e^{-b_k \Delta\tau \hat{H}_{\text{odd}}} + \mathcal{O}(\Delta\tau^5) \quad (\text{A1})$$

with suitable coefficients  $a_k$ , and  $b_k$ . In this expression,  $\hat{H}_{\text{even}}$  and  $\hat{H}_{\text{odd}}$  contain all Hamiltonian terms on odd and even bonds of the lattice, respectively, such that  $\hat{H} = \hat{H}_{\text{even}} + \hat{H}_{\text{odd}}$ . As all bond terms contained in  $\hat{H}_{\text{even}}$  are mutually commuting, it is simple to apply the unitaries  $e^{-a_k \Delta\tau \hat{H}_{\text{even}}}$  (analogously for  $\hat{H}_{\text{odd}}$ ) to matrix product states (MPS) which have the form

$$|\psi\rangle = \sum_{\{\sigma_i\}} \left( \prod_i A_i^{\sigma_i} \right) |\sigma\rangle. \quad (\text{A2})$$

Here,  $i$  runs over all lattice sites, and  $A_i^{\sigma_i}$  is a  $D_i \times D_{i+1}$  matrix. The dimensions  $D_i$  are the so-called bond dimensions and the bond dimensions for the left and right

ends of the chain are one such that the matrix product  $(\prod_i A_i^{\sigma_i})$  evaluates to a scalar.

When a time-step evolution operator such as  $\exp(-a_k \Delta\tau \hat{H}_{\text{even}})$  is applied to the MPS, the bond dimensions  $D_i$  increase to  $\tilde{D}_i \geq D_i$  – partly for technical reasons, because the resulting state needs to be brought to MPS form, and partly because the entanglement in the state can grow. States with higher entanglement generally require larger bond dimensions. Hence, it is necessary to compress the matrix product representation, i.e., to make a controlled approximation such that the bond dimensions are again reduced to some extent. This can be achieved through a Schmidt decomposition  $|\psi\rangle = \sum_{j=1}^{\tilde{D}} \lambda_j |j\rangle_L \otimes |j\rangle_R$  of the state [9] which boils down to doing singular value decompositions of the tensors  $A_i$ . The corresponding reduced density matrices for the left and right parts of the system are  $\sum_j \lambda_j^2 |j\rangle_L \langle j|_L$  and  $\sum_j \lambda_j^2 |j\rangle_R \langle j|_R$ , respectively. The bond dimension for the given bipartition of the lattice is then reduced from  $\tilde{D}$  to some value  $D' < \tilde{D}$  by retaining only the  $D'$  largest Schmidt coefficients  $\lambda_j$  and truncating all smaller ones.

$$|\psi_{\text{trunc}}\rangle = \sum_{j=1}^{D'} \lambda_j |j\rangle_L \otimes |j\rangle_R \quad (\text{A3})$$

The resulting two-norm error is

$$\left( \frac{\|\psi_{\text{trunc}} - \psi\|}{\|\psi\|} \right)^2 = \frac{\sum_{j>D'} \lambda_j^2}{\sum_j \lambda_j^2}. \quad (\text{A4})$$

In our simulations, we control errors due to truncations by discarding in every step of the algorithm only components with density matrix eigenvalues  $\lambda_j^2$  below a predefined truncation threshold. In this way, the bond dimensions become functions of (inverse) temperature and time,  $D_i = D_i(\beta, t)$ . For purifications, the truncation threshold is denoted by  $\epsilon_P$  and by  $\epsilon_M$  for the METTS. The employed values are specified in the results sections and corresponding figures. For all imaginary- and real-time evolutions of purifications, we choose time steps of size  $\Delta\tau_P = \Delta t_P = 0.125$  and for the METTS they are  $\Delta\tau_M = \Delta t_M = 0.05$ .

## Appendix B: Choices for the METTS collapse

In section III A, we reviewed the METTS algorithm for a fixed orthonormal basis  $\{|\mathbf{n}\rangle\}$ , used in the projective measurements to generate a new product state  $|\mathbf{n}'\rangle$  from the current METTS state  $|\phi_{\mathbf{n}}\rangle \equiv e^{-\beta \hat{H}/2} |\mathbf{n}\rangle / \sqrt{P_{\mathbf{n}}}$ . In order to ensure ergodicity and reduce autocorrelation times, it is useful to switch between different measurement bases  $\{|\mathbf{n}^{(k)}\rangle\}$  during the sampling [24]. This is possible, because it simply corresponds to using an overcomplete basis in the representation of the trace in Eq. (7),




Figure 10: Fit to measured computation times  $t$  for a single local tDMRG step as a function of the bond dimension  $D$ . The measured times are normalized such that  $\lim_{D \rightarrow \infty} t/D^3 \rightarrow 1$ . We fit the cost model  $c(D)$  of Eq. (C1) and obtain coefficients  $a = 7.6$ ,  $b = 17.6$ .

i.e., to starting instead from the expression

$$\langle \hat{O} \rangle_\beta = \frac{1}{Z_\beta} \sum_k \pi_k \sum_{\mathbf{n}} \langle \mathbf{n}^{(k)} | e^{-\beta \hat{H}} \hat{O} | \mathbf{n}^{(k)} \rangle, \quad (\text{B1})$$

where  $\pi_k \geq 0$  are probabilities ( $\sum_k \pi_k = 1$ ) according to which we choose basis  $\{|\mathbf{n}^{(k)}\rangle\}$  to collapse a METTS. The rest of the derivation remains unchanged except for extending the detailed balance to the choice of the measurement bases.

In our simulations, we found the following procedure to be robust and efficient. Before every collapse of the wavefunction, we randomly determine measurement bases  $\{|n_i^{(k)}\rangle, n_i = 1, \dots, d\}$  for the site Hilbert spaces  $\text{span}\{|\sigma\rangle, \sigma = 1, \dots, d\}$  and then choose  $|\mathbf{n}^{(k)}\rangle = \bigotimes_i |n_i^{(k)}\rangle$ . For every site, the local basis states  $|n_i^{(k)}\rangle$  are obtained by first choosing  $|\tilde{n}_i^{(k)}\rangle = \sum_{\sigma=1}^d (a_{n,\sigma} + i b_{n,\sigma}) |\sigma\rangle$  with  $a_{n,\sigma}$  and  $b_{n,\sigma}$  drawn from the standard normal distribution. Using the Gram-Schmidt process, these  $d$  states are subsequently orthonormalized to obtain  $\{|n_i^{(k)}\rangle\}$ .

## Appendix C: Quantification of computation costs

The main contribution of this work is to compare the efficiency of the METTS and matrix product purification algorithms. In order to make this comparison as independent as possible from specific details of the implementation and computer architecture, we approximate the computation costs associated with the imaginary- and real-time tDMRG evolutions by integrals over a function of the bond dimensions  $D_i = D_i(\beta, t)$  which were described in detail in appendix A. As the elementary steps in the employed DMRG algorithms are matrix multiplications and singular value decompositions for the MPS matrices  $A_i^{\sigma_i}$  [Eq. (A2)], the computation cost scales in leading order with the third power of the bond dimensions. How-






Figure 11: Computation time (normalized) versus theoretical computation cost according to Eq. (C2). For the XXZ model (11) with  $L = 32$  sites, we show data points for various truncation thresholds ( $\epsilon_M = 10^{-4}, 10^{-6}, 10^{-8}, 10^{-10}, 10^{-12}$ ); in each case, we include both the computation of static observables and response functions ( $t = 2.5, 5, 10$ ).

ever, small bond dimensions can lead to substantial subleading contributions corresponding to copying matrices etc. In a realistic assessment of the actual computation cost, these need to be taken into account, especially as optimal METTS bond dimensions can be quite small in some cases. Therefore, we model the computation cost  $c(D)$ , associated with a single local tDMRG step as a function of the bond dimension  $D$ , by

$$c(D) = D^3 + aD^2 + bD. \quad (C1)$$

As shown in Fig. 10, a fit to actual measured computation times yields coefficients  $a = 7.6$  and  $b = 17.6$ . While the exact values of  $a$  and  $b$  depend to some extent on the particular implementation at hand and the employed computer architecture, deviations of the coefficients do not decisively change the results of our study. In essence, the coefficients only depend on the cost scaling of the employed linear algebra routines that are implemented in the highly optimized LAPACK library [39] for the given computer architecture.

Based on this, we can model the computation costs for an imaginary-time evolution from  $\tau = 0$  to  $\beta/2$  and a subsequent real-time evolution to time  $t$ , as they occur in the different algorithms (see sections II, III A, and V), by

$$C_{\text{imag}}(\beta) = \int_0^{\beta/2} d\tau \sum_i c(D_i(\tau, 0)) \quad (C2a)$$

$$C_{\text{real}}(\beta, t) = \int_0^t dt' \sum_i c(D_i(\beta/2, t')), \quad (C2b)$$

where  $i$  runs over all bonds of the lattice. For schemes that require a certain sequence of imaginary- and real-time evolutions, we sum the corresponding contributions to quantify the total cost. For METTS, we use the aver-

age computation cost per METTS and multiply for comparisons by the considered number of METTS.

Fig. 11 shows that the theoretical computation cost (C2) provides a realistic description of the DMRG computation cost for all simulations, independent of the specific system parameters. Over several orders of magnitude, the agreement is very good; only for extremely short computation times, one obtains slight deviations. This additional overhead is mainly due to the computation time needed for the collapse procedure and the evaluation of observables.

#### Appendix D: Exact versus quasi-exact reference data

In our assessments of efficiencies, we need to quantify the errors of observables and for this purpose, reliable exact or quasi-exact data is indispensable. For the non-interacting point  $\Delta = 0$  of the XXZ model (11), the XX model, we can evaluate all observables exactly. As the system is for this case in a Gaussian state, all correlation functions are determined by the single-particle Green's function due to Wick's theorem [40]. Exploiting this, we can obtain static correlations as well as response functions by evaluating Pfaffian determinants of matrices that contain elements of the single-particle Green's function [41–43]. The colored curves in Fig. 12, marked with triangles and circles, show the errors of matrix-product purification simulations for different truncation thresholds  $\epsilon_P$  as compared to such exact results.

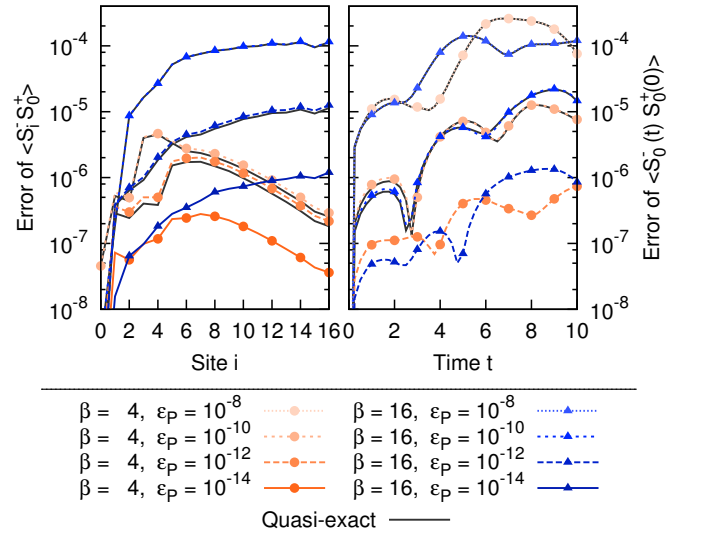


Figure 12: Deviations of purification results for different truncation thresholds  $\epsilon_P$  from exact as well as quasi-exact reference data in an XX chain [ $\Delta = 0$  in Eq. (11)] of length  $L = 32$  at inverse temperatures  $\beta = 4$  and  $16$ . The static correlator  $\langle \hat{S}_i^- \hat{S}_0^+ \rangle_\beta$  (left) and the response function  $\langle \hat{S}_0^-(t) \hat{S}_0^+(0) \rangle_\beta$  (right) are compared against an exact solution (colored curve with symbols) as well as against quasi-exact purification data (corresponding black solid curves).

Table II: Sample set sizes and bond dimensions for the comparison of static observables in the XXZ chain of length  $L = 64$  (Sec. IV B).  $D_P$  is the maximum bond dimension of the matrix product purification and  $D_M$  labels the average METTS bond dimension in the center of the chain.  $N$  denotes the sample set size that was used in the comparison of the methods.

$\beta$	$-\log_{10}(\epsilon_P)$	$-\log_{10}(\epsilon_M)$	$\Delta = 0$			$\Delta = 1$			$\Delta = 3$		
			$D_P$	$D_M$	$N$	$D_P$	$D_M$	$N$	$D_P$	$D_M$	$N$
4	10	4	22	1.0	186	29	1.0	344	42	1.0	458
4	10	6	22	2.0	54	29	2.1	80	42	2.4	183
4	10	8	22	3.8	23	29	4.5	38	42	4.8	73
4	10	10	22	5.5	13	29	7.9	16	42	7.8	30
4	12	4	28	1.0	421	51	1.0	1327	73	1.0	1779
4	12	6	28	2.0	122	51	2.1	311	73	2.4	712
4	12	8	28	3.8	53	51	4.5	147	73	4.8	286
4	12	10	28	5.5	30	51	7.9	61	73	7.8	120
16	10	6	61	2.0	467	108	2.2	1520	72	2.4	898
16	10	8	61	5.2	141	108	7.0	324	72	4.7	301
16	10	10	61	10.2	44	108	14.5	90	72	7.9	102
16	10	12	61	15.8	16	108	23.7	31	72	12.5	37
16	12	6	110	2.0	2187	204	2.2	7733	126	2.4	4735
16	12	8	110	5.2	661	204	7.0	1647	126	4.7	1590
16	12	10	110	10.2	207	204	14.5	458	126	7.9	540
16	12	12	110	15.8	79	204	23.7	158	126	12.5	196

Table III: Sample set sizes and bond dimensions for the comparison of static observables in the Bose-Hubbard model with size  $L = 32$  (Sec. IV C).  $D_P$ ,  $D_M$ , and  $N$  are defined as in Table II.

$\beta$	$-\log_{10}(\epsilon_P)$	$-\log_{10}(\epsilon_M)$	$J/U = 0.09375$			$J/U = 0.125$			$J/U = 0.25$		
			$D_P$	$D_M$	$N$	$D_P$	$D_M$	$N$	$D_P$	$D_M$	$N$
4	10	4	16	1.0	166	21	1.0	236	44	1.0	1363
4	10	6	16	2.1	52	21	2.6	52	44	3.0	239
4	10	8	16	3.3	26	21	3.7	31	44	5.4	92
4	10	10	16	5.5	10	21	6.3	12	44	9.8	29
4	12	4	31	1.0	588	40	1.0	1128	66	1.0	4684
4	12	6	31	2.1	186	40	2.6	251	66	3.0	821
4	12	8	31	3.3	93	40	3.7	148	66	5.4	317
4	12	10	31	5.5	38	40	6.3	59	66	9.8	102
16	10	6	33	2.1	246	44	2.7	324	99	3.1	2019
16	10	8	33	3.2	119	44	3.8	176	99	6.5	543
16	10	10	33	7.0	33	44	8.5	42	99	11.9	142
16	12	6	57	2.1	1008	81	2.7	1651	187	3.1	12204
16	12	8	57	3.2	487	81	3.8	898	187	6.5	3285
16	12	10	57	7.0	138	81	8.5	215	187	11.9	864

Although the XXZ model (11) is Bethe-Ansatz integrable [28] for all values of  $\Delta$ , no exact expressions for longer-ranged static correlators or response functions are available. The same is true for the Bose-Hubbard model (12) which is not (known to be) integrable. For these cases, we use very accurate purification simulations to obtain quasi-exact reference data. In particular, we used the truncation threshold  $\epsilon_P = 10^{-14}$  during imaginary-

time evolutions and  $\epsilon_P = 10^{-12}$  during (subsequent) real-time evolutions. ( $\epsilon_P$  has been defined in appendix A.) To justify this procedure, we show in Fig. 12 with black solid lines (no symbols) also the errors of the other simulations compared to this quasi-exact data. These curves deviate indeed only very slightly from the colored curves that correspond to the comparison to the exact results.

- 
- [1] S. R. White, Phys. Rev. Lett. **69**, 2863 (1992).
  - [2] S. R. White, Phys. Rev. B **48**, 10345 (1993).
  - [3] U. Schollwöck, Rev. Mod. Phys. **77**, 259 (2005).
  - [4] G. Vidal, Phys. Rev. Lett. **93**, 040502 (2004).
  - [5] S. R. White and A. E. Feiguin, Phys. Rev. Lett. **93**, 076401 (2004).
  - [6] A. Daley, C. Kollath, U. Schollwöck, and G. Vidal, J. Stat. Mech. P04005 (2004).
  - [7] A. Uhlmann, Rep. Math. Phys. **9**, 273 (1976).
  - [8] A. Uhlmann, Rep. Math. Phys. **24**, 229 (1986).
  - [9] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).
  - [10] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, Phys. Rev. Lett. **93**, 207204 (2004).
  - [11] M. Zwolak and G. Vidal, Phys. Rev. Lett. **93**, 207205 (2004).
  - [12] A. E. Feiguin and S. R. White, Phys. Rev. B **72**, 220401(R) (2005).
  - [13] T. Barthel, diploma thesis, RWTH Aachen, 2005.
  - [14] T. Barthel, U. Schollwöck, and S. R. White, Phys. Rev. B **79**, 245101 (2009).
  - [15] A. E. Feiguin and G. A. Fiete, Phys. Rev. B **81**, 075108 (2010).
  - [16] C. Karrasch, J. H. Bardarson, and J. E. Moore, Phys. Rev. Lett. **108**, 227206 (2012).
  - [17] T. Barthel, U. Schollwöck, and S. Sachdev, arXiv:1212.3570 (2012).
  - [18] T. Barthel, New J. Phys. **15**, 073010 (2013).
  - [19] C. Karrasch, J. H. Bardarson, and J. E. Moore, New J. Phys. **15**, 083031 (2013).
  - [20] C. Karrasch, J. Hauschild, S. Langer, and F. Heidrich-Meisner, Phys. Rev. B **87**, 245128 (2013).
  - [21] B. Lake, D. A. Tennant, J.-S. Caux, T. Barthel, U. Schollwöck, S. E. Nagler, and C. D. Frost, Phys. Rev. Lett. **111**, 137205 (2013).
  - [22] Y. Huang, C. Karrasch, and J. E. Moore, Phys. Rev. B **88**, 115126 (2013).
  - [23] S. R. White, Phys. Rev. Lett. **102**, 190601 (2009).
  - [24] E. M. Stoudenmire and S. R. White, New J. Phys. **12**, 055026 (2010).
  - [25] Y. Yao, Y. Zhan, X. Hou, and C.-Q. Wu, arXiv:1206.4372 (2012).
  - [26] G. Alvarez, Phys. Rev. B **87**, 245130 (2013).
  - [27] L. Bonnes, F. H. L. Essler, and A. M. Läuchli, arXiv:1404.4062 (2014).
  - [28] H. A. Bethe, Z. Phys. **71**, 205 (1931).
  - [29] J. D. Cloizeaux and M. Gaudin, J. Math. Phys. **7**, 1384 (1966).
  - [30] H.-J. Mikeska and A. K. Kolezhuk, in *Quantum Magnetism*, Vol. 645 of *Lecture Notes in Physics*, edited by U. Schollwöck, J. Richter, D. J. J. Farnell, and R. F. Bishop (Springer, Berlin, 2004), pp. 1–83.
  - [31] T. D. Kühner and H. Monien, Phys. Rev. B **58**, R14741 (1998).
  - [32] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller, Phys. Rev. Lett. **81**, 3108 (1998).
  - [33] U. Schollwöck, Physics **2**, 39 (2009).
  - [34] I. Pizorn, V. Eisler, S. Andergassen, and M. Troyer, New J. Phys. **16**, 073007 (2014).
  - [35] B. Nachtergaele and R. Sims, in *New Trends in Mathematical Physics. Selected contributions of the XVth International Congress on Mathematical Physics*, edited by V. Sidoravicius (Springer, Heidelberg, 2009), pp. 591–614.
  - [36] T. Barthel and M. Kliesch, Phys. Rev. Lett. **108**, 230504 (2012).
  - [37] H. F. Trotter, Proc. Am. Math. Soc. **10**, 545 (1959).
  - [38] M. Suzuki, Prog. Theor. Phys. **56**, 1454 (1976).
  - [39] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen, *LAPACK Users' Guide*, 3rd ed. (Society for Industrial and Applied Mathematics, Philadelphia, PA, 1999).
  - [40] A. L. Fetter and J. D. Walecka, *Quantum theory of many-particle systems* (McGraw-Hill, San Francisco, 1971).
  - [41] E. R. Caianiello and S. Fubini, Nuovo Cimento **9**, 1218 (1952).
  - [42] H. S. Green and C. A. Hurst, *Order-disorder phenomena*, Vol. 5 of *Monographs in statistical physics* (Interscience Publishers, London, 1964).
  - [43] J. Stolze, A. Nöppert, and G. Müller, Phys. Rev. B **52**, 4319 (1995).
  - [44] A state  $|\mathcal{P}_{\hat{\rho}}\rangle \in \mathcal{H} \otimes \mathcal{H}_{\text{aux}}$  is called a purification of the density matrix  $\hat{\rho}$  on  $\mathcal{H}$  if  $\text{Tr}_{\text{aux}}|\mathcal{P}_{\hat{\rho}}\rangle\langle\mathcal{P}_{\hat{\rho}}| = \hat{\rho}$ .
  - [45] Computation costs for DMRG simulations with open boundary conditions scale as  $D^3$  with the bond dimension  $D$ .
  - [46] For operators  $\hat{Y} = \hat{X}^\dagger$  and Hamiltonians which are real in the chosen onsite basis, the distribution of individual METTS measurement results in schemes A and B will be the same, as the observable  $\hat{X}(t)\hat{X}^\dagger(0)$ , evaluated in scheme A, is then just the transpose of  $\hat{X}(0)\hat{X}^\dagger(-t)$ , evaluated in scheme B. In general, the distribution of individual measurement results and hence the statistical errors of schemes A, B, and C can differ.